

# METHOD AND SYSTEM FOR AUTOMATICALLY PRESERVING PERSISTENT STORAGE

## CROSS-REFERENCE TO RELATED APPLICATION

This application claims the benefit of U.S. Provisional Patent Application  
5 No. 60/223,829 filed August 8, 2000, which is incorporated herein by reference in its  
entirety.

## BACKGROUND OF THE INVENTION

### Field of the Invention

10 The present invention relates to a method and system for protecting a  
computer system, and, in particular, to methods and systems for providing a storage  
redirection driver that protects the storage devices of a computer system from alteration.

### Background

Often times it is desirable to insure that a computer workstation remains  
15 unaltered, even though it is being used by a multitude of users in potentially varied and  
unknown ways. For example, in public environments such as schools, libraries, and other  
community facilities, access to a group of computer systems is highly desired, for example,  
to conduct research, use common resources of a company, or to learn and try new  
programs. Typically, these environments can be characterized as having “uncontrolled  
20 access,” because the moment by moment usage of the computer systems is not monitored.

In such situations, it is highly desirable to insure that users can utilize the  
functionality of the computer system but not be allowed to “corrupt” the persistent storage  
of the system, so that follow-on users will find the system in a pristine state. One  
mechanism for preventing such corruption is to make a copy of the pristine state of the  
25 computer system desired and to restore the copy prior to shutdown of the system by each

user. One difficulty with such an approach is that it is often difficult in operating system environments to capture the hardware on/off switch, and so, when a user cycles power without a proper operating system shutdown, the system is unable to successfully restore the pristine state.

5           Some systems have used special hardware cards to perform the copying. For example, one such system copies the pristine state of the disk storage device into a special disk partition before use access is permitted. Data from the special disk partition is then copied back (restored) upon computer system shut down.

10           In database environments, other techniques have been used. Typically, for example, the database changes (in opposite order) are “undone” to the storage device to return it to a pristine state. This technique also suffers from failure in the cycle power situation.

## SUMMARY OF THE INVENTION

15           Embodiments of the present invention provide methods and systems for automatically preserving an original state of a computer system upon rebooting. Example embodiments provide a Disk Access Redirection System (the “DARS”) to allow all or portions of a storage device to be protected from modification. The DARS can protect such storage devices as disk drives, and other persistent and semi-persistent storage devices. The DARS reads data from and writes data to a redirected data area (a redirected  
20   space) when a storage access request is received that would otherwise alter the state of an area of the storage device that has been designated as protected. When the computer system is shut down, the redirected data area is discarded; thus, when the computer system is rebooted, the original state of the protected portions of the storage device are automatically restored, without the need to copy any information from a backup area.

25           In one embodiment, the DARS comprises a Redirection Driver, and several redirection tables, including an Available Space Table (“AST”), a Protected Space Redirection Table (“PSRT”), and, when unprotected areas can be designated, an Unprotected Space Table (“UST”). The AST indicates available space of a storage device

and is used to allocate space as redirected space. The PSRT indicates the mappings of protected locations on the storage device to locations in the redirected space. The UST indicates unprotected locations of the storage device. In some embodiments, one or more entire storage devices can be designated as protected. In other embodiments, portions of a storage device can be designated as protected and other portions designated as unprotected. The DARS differentiates between protected and unprotected storage areas, redirecting storage write requests to a redirected data area when a storage area has been designated as protected.

In one embodiment, the protection level for the DARS is configurable. For example, the protection level may be configured as all storage is protected, portions of storage (at the device level or within a device) are protected, or no storage is protected. Variations of these combinations are also configurable.

In some embodiments, when a portion of storage is indicated as unprotected, it is written to directly by the appropriate storage driver. In other embodiments, the Redirection Driver redirects storage access requests to unprotected areas, and integrates them into the unprotected portion of storage upon computer system shut down, or upon reboot.

In one embodiment, the Redirection Driver translates an original location in a storage access request to a redirected storage location and forwards a revised storage access request to the storage driver to perform the storage access. In another embodiment, the Redirection Driver performs the actual access itself.

In yet another embodiment, the Redirection Driver can redirect data at different driver access levels. For example, the Redirection Drive can direct data at a file level, a cluster level, a logical sector level, or a physical sector level. Combinations of redirecting data at different levels are also provided. In addition, a Redirection Driver that operates using virtual clusters (or any data abstraction that can be implemented as any size object for storing data) can be used with the DARS. In this embodiment, available space objects are provided that implement the mapping between virtual clusters and the actual

size of the space in the redirected area. Virtual cluster embodiments provide extensibility across operating systems and different storage devices.

In one embodiment that supports a layered driver architecture, the Redirection Driver is inserted into a chain of native operating system drivers and registered with the operating system. When a storage access request is made, the appropriate driver in the chain is invoked by the operating system. In this manner, the Redirection Driver intercepts storage access requests. In some such embodiments, the layer at which the Redirection Driver is inserted into the chain of drivers is based upon the level of storage access being redirected.

In some embodiments, the DARS optionally performs functions upon computer system shutdown. In one such embodiment, the DARS optionally saves the AST, PSRT, and UST tables so that the redirected data can be restored upon computer system reboot. In another embodiment, redirected data that corresponds to unprotected areas of storage is also saved.

In one embodiment, the DARS saves the AST and PSRT tables at a desired frequency to the hard disk, so that redirected data can be restored upon a system reboot.

## BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a graphic illustration of an example use of a Disk Access Redirection System (DARS) in a library.

Figure 2 is an example block diagram of disk access redirection performed by the Disk Access Redirection System.

Figure 3 is a block diagram of the components of an example embodiment of a Disk Access Redirection System.

Figure 4 is an overview flow diagram of the example operation of the Disk Access Redirection System to achieve storage preservation according to techniques of the present invention.

Figure 5 is an example flow diagram of the steps performed by the Disk Access Redirection System when the Redirection Driver intercepts a disk access request.

Figure 6 is an example block diagram of a general purpose computer system for practicing preferred embodiments of the present invention.

Figure 7 is an example block diagram of a typical layout of a physical disk storage device in a computer system.

5           Figure 8 is an example block diagram of a disk drive partitioned into logical drives (partitions).

Figure 9 is an example block diagram of a layered I/O driver architecture of a typical operating system with an inserted Redirection Driver.

Figure 10 is an example block diagram illustrating how the Disk Access  
10   Redirection System tables are used by the Redirection Driver to redirect disk accesses.

Figure 11 is an example block diagram of a disk drive with a File Allocation Table arranged according to the FAT 16 architecture.

Figure 12 is an example block diagram of a disk drive with a File Allocation Table arranged according to the FAT 32 architecture.

15           Figure 13 is an example block diagram of a disk drive arranged according to the New Technology File System ("NTFS") architecture.

Figure 14 is an example flow diagram of the steps performed by the Disk Access Redirection System when the computer system boots up.

Figure 15 is an example flow diagram of the steps performed by the Disk  
20   Access Redirection System when a new device is mounted by the computer system.

Figure 16 is an example flow diagram of the steps performed by the Redirection Driver when a read request is intercepted.

Figure 17 is an example flow diagram of the steps performed by the Redirection Driver when a write request is intercepted.

25           Figure 18 is an example flow diagram of the steps performed by the Redirection Driver when a system shutdown request is intercepted.

## DETAILED DESCRIPTION OF THE INVENTION

Embodiments of the present invention provide methods and systems for preserving an original state of a computer system upon rebooting. Common operating systems and application programs rely upon storing data to and retrieving data from persistent storage devices such as hard disks and other types of memory. Example embodiments of the present invention provide a Disk Access Redirection System (the “DARS”) to allow all or portions of such storage to be protected from modification. Using the DARS, operating systems, application programs, and other code read and write data on a computer system storage device in a way that automatically allows the original state of protected portions of the storage device to be re-established when the computer system is rebooted. Specifically, the DARS evaluates storage access requests and automatically directs or redirects the flow of information so that, transparent to the user, the original state of the protected storage device is maintained and the changes to protected portions of the storage device are discarded when the computer system is rebooted. In this manner, security of the storage device is insured – even from a user who powers off the machine using a power switch without properly shutting down the system. Although primarily discussed below with reference to disk storage and access, especially hard disks, one skilled in the art will recognize that the techniques of the present invention are also applicable to other types of persistent and semi-persistent storage devices, including such devices as: CDROMS, Flash Memory, floppy disks, and other types of removable media storage devices.

The Disk Access Redirection System is useful in a multitude of situations, including those in which it is desirable to allow operators (users) of the computer system to actually use the system to produce useful output, but where the risk of leaving the system in an inoperable state is too high. For example, the DARS may be useful in a teaching situation where it is highly desirable that the computer system isn’t truly modified after each teaching session and is returned to its initial state before each session. Or, for example, the DARS could be used in a library where a variety of persons with different levels of experience may need to use computer systems to locate particular pieces of

literature, perform database searches, or edit documents (public access word processing, for example). In these situations, it is desired that the machine can be easily (and speedily) returned to some known, operable state.

Figure 1 is a graphic illustration of an example use of a Disk Access  
5 Redirection System (DARS) in a library. A librarian 101 operates and maintains a  
computer system 103 for the benefit of various users 102. The computer system 103  
includes a storage device 104 (for example, a hard disk) with available space 108. The  
users 102 may want to use the computer system 103, for example, to access an online card  
10 catalog, access the Internet, or perhaps even to download and execute programs on the  
computer system 103. Without the use of the DARS, the users 102 may store files on the  
hard disk 104 that may quickly exhaust the available space 108 or may execute malicious  
code, for example, that installs viruses on the hard disk 104. The DARS allows the  
librarian 101 to easily restore the original state of the disk, which preferably contains  
sufficient available space 108 and does not contain malicious code, by simply rebooting the  
15 system.

Specifically, the librarian 101 first installs and activates the DARS on the  
computer system 103. The DARS identifies space on the disk that is to be protected,  
protected space 106, and space that is available, available space 108, (e.g., space currently  
unused by any programs). Optionally, if only a portion of the disk is to be protected,  
20 unprotected space 105 may also be identified. As needed, the DARS allocates available  
space 108 to a "redirected" space 107. When the users 102 execute programs that request  
data to be written to the protected space 106 on the disk, the DARS automatically  
intercepts the write request and redirects the write to memory in the redirected space 107  
on the disk. When the users 102 execute programs that request data to be read from the  
25 disk, the DARS intercepts the read request and, automatically determines from which  
location the read should be performed. When the librarian 101 or one of the users 102  
shuts down the computer system 103, data stored in the redirected space 107 is discarded,  
rendering any apparent changes to the protected space 106 "lost" from the user's  
perspective and "ignored" from the system's perspective; when the computer system 103 is

rebooted, the protected space 106 is unchanged from the state it occupied prior to installation of the DARS. In one embodiment, the DARS supports the ability to protect some but not all portions of the disk. In this case, when the computer system 103 is rebooted, modifications written to the unprotected space 105 remain.

5           Figure 2 is an example block diagram of disk access redirection performed by the Disk Access Redirection System. In summary, the DARS operates by intercepting disk access requests before the appropriate device driver handles the disk access requests. A disk access request 201 is formed, for example, by an application program or an operating system function, and includes indicators that identify the type of request and the  
10   memory address to which the access is being requested. The DARS 202 intercepts the disk access request 201 before it is sent to the appropriate disk driver 203. The disk driver 203 then reads from or writes to the designated area on disk drive 204.

          In a typical operating system, storage devices are accessed by device drivers that are organized according to a layered architecture. For example, in the Windows 9X or  
15   NT systems, developed by Microsoft Corporation, application requests to access a file are passed to one or more high level drivers, which process and eventually forward the request into one or more requests to a low level driver that communicates with a device using physical address indicators. The DARS 202 can be implemented to intercept disk access requests at several of these levels. Which driver(s) is (are) intercepted by the DARS 202  
20   depends upon whether redirection is occurring on the file, cluster, or sector level, or according to some other logical drive division scheme implemented by the computer system or by the DARS 202.

          Once the request is intercepted, the DARS 202 evaluates the type of the request and determines whether the requested original address or a redirected address  
25   should be accessed. The DARS 202 then forwards the disk access request 201 with the original address or a redirected address to the appropriate drivers 203. The disk drivers 203 then handle the disk access request 201, accessing as indicated, the protected space 206, the unprotected space 205, or the redirected space 207. One skilled in the art will recognize that in an alternative embodiment, the DARS 202 may handle the disk access request itself



rather than forwarding the request to native operating system disk drivers. Other combinations, such as forwarding the request when the original address is used and handling it otherwise, are also contemplated.

Figure 3 is a block diagram of the components of an example embodiment of a Disk Access Redirection System. The example embodiment of a Disk Access Redirection System 301 comprises a Redirection Driver 302, an Available Space Table (AST) 303, a Protected Space Redirection Table (PSRT) 304, and, optionally, an Unprotected Space Table (UST) 305. One skilled in the art will recognize that the AST 303, PSRT 304, and UST 305 may be stored in any type of storage, including for example, volatile memory. The Redirection Driver 302 intercepts the disk access request, determines whether or not the requested memory address is to be redirected, and forwards the disk access request to the next appropriate Input/Output ("I/O") driver. (See, for example, Figure 2.) The AST 303 is used by the DARS 301 to map available memory for use as the redirected space (*e.g.*, available space 208 and redirected space 207 in Figure 2). The PSRT 304 is used by the DARS 301 to manage the mappings of protected space to redirected space (*e.g.*, protected space 206 in Figure 2). The UST 305 is used by the DARS 301 to map the unprotected space so that it is either written to directly or redirected to space that is preserved across boot sessions. The UST 305 is optional in that there may be implementations where all storage space is either protected or available, and none of the storage space is allowed to be unprotected (alterable).

Figure 4 is an overview flow diagram of the example operation of the Disk Access Redirection System to achieve storage preservation according to techniques of the present invention. In step 401, a user turns on the computer system. In step 402, the computer system loads the Redirection Driver. For example, in a computer system running the Microsoft Windows 9X Operating System, the redirection driver is stored in a special folder, for example the "windows\system\iosubsys" folder to be automatically loaded when the operating system boots up. In addition, the driver indicates to the operating system at which location in the layered architecture it desires to be installed. In another example, a computer system running the Microsoft Windows NT or 2000 Operating System, the

redirection driver is registered with the operating system during an installation procedure. When the computer is booted, the redirection driver is loaded by the operating system. One skilled in the art will recognize that the specific steps necessary to install, register, initialize and/or load the redirection driver may vary depending on requirements of the  
5 operating system or other system components. These system dependent modifications are contemplated and are to be included within the scope of the invention.

In step 403, the computer system determines the desired protection level. For example, the computer system may allow a user, for example, a system administrator or other operator, to access a password protected user interface, through which the user  
10 specifies a desired level of system protection. The user interface may be implemented, for example, as a portion of the DARS redirection driver (which grabs control of the system right after it boots up) or as a separate part of the operating system installed as an initial start-up program. Example protection levels may include no protection, partial protection, and full protection. These protection levels correspond to the amount of security desired,  
15 where full protection insures no alteration of the persistent storage devices. In step 404, the system examines the determined protection level, and if no protection is desired, the DARS is effectively terminated in that the Redirection Driver is not initialized to intercept requests; otherwise, the system continues in step 405.

In step 405, if the DARS determines that the computer system is to be  
20 partially or fully protected, it registers the Redirection Driver to intercept disk access requests at the appropriate level as discussed with reference to Figure 2. For the purposes herein, "registered" implies any actions (including none) that are required by a particular operating system to make the Redirection Driver known to the operating system so that the Redirection Driver can intercept access requests. In some cases, "registered" means that  
25 the Redirection Driver calls a particular operating system function to become known. In other cases, such as within Windows 9X, the driver is automatically registered to intercept requests at a location in the driver chain according to data stored in the Redirection Driver file at bootup time provided the driver is stored in the "windows\system\iosubsys" directory. In systems running the Windows NT operating system, data stored in the

operating system registry when the driver is installed is used to register the driver at the correct level each time the system is booted.

In step 406, as disk access requests are received, the Redirection Driver intercepts the disk access requests. The steps performed by the DARS when the Redirection Driver intercepts disk access requests are discussed below in detail, with respect to Figure 5. Typically, at some later point in time, in step 409, the computer system shuts down. Optionally, steps 407 and 408 are performed before the computer system shuts down. In step 407 the Redirection Driver intercepts the system shutdown request. In step 408, based on the determined protection level, the DARS may either save modifications that were made or discard them.

Figure 5 is an example flow diagram of the steps performed by the Disk Access Redirection System when the Redirection Driver intercepts a disk access request. As discussed above, with respect to Figure 4, the computer system determines the protection level for the system. If no protection level is desired, the Redirection Driver is not initialized (or registered to intercept access requests), so has no effect on disk accesses. If at least partial protection of the computer system is desired, then the Redirection Driver is registered at the appropriate level in the system, for each storage device to be protected, and the following steps are performed when the Redirection Driver intercepts a disk access request.

Specifically, in step 501, the Redirection Driver determines whether the access requested is for a read or a write. If the request is for a read, the Redirection Driver continues in step 502, else, it continues in step 504. In step 502, the Redirection Driver determines whether or not the read is being requested against memory in the protected space (*e.g.*, protected space 206 in Figure 2). If the read is being requested against memory in the protected space (no alteration is allowed), the Redirection Driver continues in step 503, else it continues in step 508. In step 503, the Redirection Driver determines whether or not the requested memory address in the protected space already has been redirected to redirected space (*e.g.*, redirected space 207), and, if so, continues in step 507, else continues in step 508. In step 507, the Redirection Driver determines the redirected

address. In step 508, the Redirection Driver forwards the read request with the original or redirected address as determined to the next appropriate I/O driver.

When it is determined in step 501 that the requested disk access is for a write, then, in step 504, the Redirection Driver determines whether or not the write is being requested to memory in the protected space (the write is “unallowed”), and, if so, continues in step 505, else, continues in step 508 to forward the write request to the next appropriate I/O driver. In step 505, the Redirection Driver determines whether or not the requested memory address has been previously redirected (for example, due to a prior write request), and, if so, continues in steps 507 and 508 as described, else continues in step 506. In step 506, the Redirection Driver allocates an available memory address to be used for redirected space and maps the requested memory address to an address in the redirected space. The Redirection Driver then continues in step 508 as described above. In an alternate embodiment, in step 508, the Redirection Driver handles the disk access requests rather than forwarding them to a native operating system I/O driver.

Figure 6 is an example block diagram of a general purpose computer system for practicing preferred embodiments of the present invention. The computer system 601 contains a central processing unit (CPU) 602, input/output devices 603, including storage device 613, a display device 612, and a computer memory (memory) 604. The Disk Access Redirection System 611, comprising Redirection Driver 605, Available Space Table 608, Protected Space Redirection Table 609, and Unprotected Space Table 610 preferably resides in memory with the operating system 606 and other programs 607. The storage device 613 contains an unprotected space 614, a protected space 615, a redirected space 616, and an available space 617. One skilled in the art will recognize that the storage device 613 may be implemented in any of various configurations including, for example, one or more physical or virtual disk drives located on one computer or located on multiple computers connected through a network.

One skilled in the art also will recognize that various arrangements of this computer system and its components are possible and contemplated by the methods and systems of the present invention. For example, the various tables of the Disk Access

Redirection System may reside in separate memories or span across several memories or be non-contiguous. Well-known techniques for handling such data structures and memory management can be used. In addition, the Redirection Driver, before being installed into memory, may be remotely located and accessed for use via a network when desired.

- 5 Various other modifications to the storage organization and the location of the other parts of the computer system are also contemplated. In addition, in the example flow diagrams described, different orderings of the steps and different divisions of the steps are likewise contemplated to accomplish the techniques of the present invention.

- As previously mentioned, the Redirection Driver 605 may be implemented  
10 to intercept I/O (input/output) requests at various levels in the storage driver architecture of the operating system 606 depending upon what level of abstraction of storage is being redirected. For example, the Redirection Driver 605 may be implemented to redirect data access at the file level or at the physical sector address level, or at some other level. Additionally, multiple Redirection Drivers may be used, for example, one Redirection  
15 Driver for each storage device type, or for different levels of access.

- Figure 7 is an example block diagram of a typical layout of a physical disk storage device in a computer system. As shown on the left side of Figure 7, a disk storage device 701 is made up of a series of sectors 730. Each sector may be identified by (1) a head number (*e.g.*, 720), which defines the disk surface where the sector resides; (2) a  
20 cylinder number (*e.g.*, 740), which defines the track where the sector resides; and (3) a sector number (*e.g.*, 730) within the associated track. These identifying attributes may then be used to map a unique number to each physical sector. The right side of Figure 7 shows the same disk storage device 701, but using a contiguous sector description 702. For example, the first sector is assigned the number 00, the second is assigned the number 01, and so on. This allows each sector to be accessed through a simple indexing scheme, 0, 1,  
25 2 ... (n-1), where n is the total number of sectors on the physical disk. Each sector 702 comprises a fixed number of bytes 703. For example, Sector 00 contains 8 bytes, labeled Byte 0 – Byte 7.

In addition, a physical disk may be partitioned into one or more logical drives. Figure 8 is an example block diagram of a disk drive partitioned into logical drives (partitions). The disk drive is shown divided into 3 logical partitions: 802, 803, and 804. Within a partition, the sectors are referred to as logical sectors. For example, partition 802 contains physical sectors 812-818 which are mapped to logical sectors 0-6, and partition 803 contains physical sectors 820-823 mapped to logical sectors 0-3. Logical sectors are accessed by a simple indexing scheme 0, 1, 2, ... (n-1) relative to the partition, where n is the number of logical sectors in the partition. The size in bytes of a logical sector is not necessarily the same as a physical sector, but may be an integral number of physical sectors. For example, partition 804 contains physical sectors 840-845 mapped to logical sectors 0-2, where each logical sector is equal in size to two physical sectors. Although logical partitions 802, 803, and 804 are shown mapped to contiguous physical sectors, in some systems, logical partitions may be mapped to non-contiguous physical sectors.

User or application information and other data are typically stored in files. A file can be stored on a logical drive in one or more logical sectors, which may or may not be contiguous. Different operating systems use different schemes to keep track of where a file is stored on a logical drive and to keep track of which sectors a file occupies. Often times some sort of "table," "database," or other appropriate data structure, is used to locate the file on the logical drive and to track the sectors allocated to the file. In some systems, to reduce the size of the database required to describe the files on a disk storage device, the abstraction of a cluster is used. A cluster is an integral number of sequential sectors. For example, in the Windows 9X operating system, a cluster can be up to 64 logical sectors. The way that an operating system organizes files into clusters and the way it keeps track of this organization in a "database" is typically referred to as the file system. The database used by the file system must also be stored on the disk storage device. This database can be stored on the logical drive as a sequential set of logical sectors that is fixed in size and location, or it can be stored as a special file where access to it is achieved through a smaller fixed set of data.

In any operating system and file system, user and application data starts out in the form of a file. The operating system maintains a database of the clusters that are used to store the file data. Any read or write of a file or portion of a file ends up being a read or write of one or more physical sectors. Modern operating systems are typically  
5 coded in such a way that the file system reads and writes to a disk storage device through a layering of drivers as briefly described above. These drivers are defined such that drivers at the top communicate at the file level, while drivers at the bottom communicate at the physical sector, level of the disk storage device. At some point in the layered input/output system, a driver can be inserted to intercept all the reads and writes at the level it is  
10 inserted, for example, at either the logical sector level or the physical sector level.

Figure 9 is an example block diagram of a layered I/O driver architecture of a typical operating system with an inserted Redirection Driver. The layered I/O driver architecture 901 comprises file level drivers 902, logical sector level drivers 903, and physical sector level drivers 904. In Figure 9, the Redirection Driver 920 implements  
15 logical sector redirection, so is shown inserted before the other drivers that handle logical sectors 903. When a request for file access comes in, a file level driver 902 translates the request into a read or a write request of the appropriate logical sectors and forwards the translated request down the driver chain. The Redirection Driver 920, which has been inserted into the driver chain before logical sector drivers 903, receives the translated  
20 request and processes the request according to the present invention. One skilled in the art will recognize that the Redirection Driver 920 can be similarly inserted at other points (or at multiple points) in the driver chain, depending upon what disk access abstraction has been implemented in the driver.

In example embodiments of the Disk Access Redirection System, the  
25 Available Space Table 608, Protected Space Redirection Table 609, and Unprotected Space Table 610 may be organized at the logical sector level, at the cluster level, at some other level of data abstraction, or a combination of data abstraction levels. When organized at the sector level, the Protected Space Redirection Table 609 tracks all sectors that can be redirected, and each entry corresponds to one logical sector. For example, if an entry for a

particular sector is zero, then no redirection has occurred. In contrast, if an entry is non-zero, then redirection has occurred, and an indicator of the redirected sector is stored in that table entry. (The actual entry may contain a variety of data, such as a logical sector address in Redirected Space, or a pointer to an entry in the AST 608, which is allocated to the  
5 Redirected Space.) When organized instead at the cluster level, the Protected Space Redirection Table 609 tracks all clusters that can be redirected, and each entry in the table corresponds to one cluster. For example, if an entry for a particular cluster is zero, then no redirection has occurred. In contrast, if an entry is non-zero, then redirection has occurred, and an indicator of the number of the redirected cluster (or other reference) is stored in the  
10 table entry.

The basis of redirection may be clusters, sectors, or any other level of data abstraction, and is based on the organization of the DARS. In some embodiments, the use of clusters will reduce the memory size of the DARS. One embodiment organizes the DARS tables as groups of clusters or sectors instead of as individual clusters or sectors.  
15 When using groups of clusters, a table entry consists of a cluster number and an extent (*i.e.*, a number of sequential clusters). Because operating systems attempt to keep reads and writes of related data in sequential clusters, the first write to a cluster will usually be a write to a sequence of clusters. For example, if a write operation calls for data to be stored in clusters 5, 6, 7, 8, 9, and 10, this information may be stored in a table where each element  
20 in the table consists of 3 pieces of information: (1) the original starting cluster number, (2) the number of sequential clusters, and (3) the number of the first redirected cluster. In this example, storing the redirection information in the Protected Space Redirection Table 609 requires only three numbers, rather than the six required when redirection information is referred to as individual clusters. As the write blocks get larger, the savings in memory  
25 become significant. No matter how large the write becomes, only three numbers are required to store all the redirection information. The other DARS tables can be similarly implemented in terms of logical sectors, clusters, or a combination of both.

Figure 10 is an example block diagram illustrating how the Disk Access Redirection System tables are used by the Redirection Driver to redirect disk accesses. As



described earlier, disk storage comprises protected space 1001, which cannot be altered; unprotected space 1004, which can be altered; redirected space 1002, which stores attempted alterations to the protected space 1001; and available space 1003, which is currently unused space. The tables of the DARS (tables 1020, 1030, and 1040) are used to manage the mappings between the protected space 1001, redirected space 1002, available space 1003, and unprotected space 1004. The Available Space Table 1030 (“AST”) maps the available space 1003 (shown hatched); the Protected Space Redirection Table 1020 (“PSRT”) maps the protected space 1001 to the redirected space 1002; and the Unprotected Space Table 1040 (“UST”) maps/records the unprotected space 1004 (shown hatched).

In one embodiment, the AST 1030 stores a series of records, each record indicating a portion of memory that is available for redirection use. The AST may be organized as a simple list of available cluster numbers, or it may be organized as a table of items, each of which contain a starting cluster number and an extent (*i.e.*, number of sequential empty clusters). In the embodiment shown, each AST record comprises an available address 1031 that indicates the beginning of an available portion of memory and an extent 1032 that indicates the length of the available portion of memory (*e.g.*, the first shown record indicates two clusters of storage with addresses A01 through A02). As empty clusters are used for redirection, they are removed from the AST or marked as not available. A write operation of “n” sequential clusters generates a request for “n” sequential empty clusters for redirection. The AST services this request. If it is not possible to find “n” sequential clusters in the AST, the request is preferably broken up into smaller units.

The UST 1040 is similar to the AST 1030, storing a series of records, each comprising a beginning address 1041 and an extent 1042 (*e.g.*, the first shown record indicates two clusters of storage with addresses U01 through U02). The records in the UST indicate portions of unprotected space 1004, which can be altered. In an alternative embodiment, the unprotected space 1004 is mapped via the UST 1040 to redirected space 1002 and saved upon computer shut down. This embodiment has the disadvantage of potentially losing data written to the unprotected space 1004 when the user simply powers

down the machine using the hardware power switch, without gracefully shutting down the operating system and allowing the redirected space 1002 to be saved.

The PSRT 1020 maps portions of protected space 1001 to redirected space 1002. Each record in the PSRT 1020 comprises a protected address 1021 that indicates the beginning of a protected portion of storage, an extent 1022 that indicates the length of the protected portion of storage, and a redirected address 1023 that indicates the beginning of a portion of memory to which all reads and writes to the indicated protected address of that record will be redirected. For example, in the first record, addresses P01 through P03 of storage in the protected space 1001 (as shown by an extent of 3 clusters) is currently redirected to addresses R02 through R04 in the redirected space 1002. Although described with respect to clusters, one skilled in the art will recognize that a similar organization and description can be used to implement the tables and the space layout using other abstractions, such as at the file level or sector level.

In an example embodiment shown, the AST, UST, and PSRT (*e.g.*, AST 1030, UST 1040, and PSRT 1020 in Figure 10) are stored in a volatile memory of the computer system. In alternate embodiments, these tables may be written to persistent storage and deleted prior to the computer system being shut down. In another alternate embodiment, the determined system protection level may indicate whether the tables are stored in volatile memory or persistent storage. For example, one level of protection may be defined which maintains the AST, PSRT, and UST in persistent storage across several system reboots. This may be beneficial on a system that is being used to test new software. When testing is complete, or at various stages during the testing process, the system protection level could be modified to a level that causes the DARS to store the AST, PSRT, and UST in volatile memory or to delete them from persistent memory on shutdown. As one skilled in the art will recognize, various protection levels may be defined, each requiring possible variations in implementation of the DARS system tables. These variations may include the specific data structures of the tables in addition to the type of memory used to store the tables. Well-known techniques for varying the data structures

and for indicating mappings can be used. All of these variations are contemplated and are intended to fall within the scope of the invention.

In some operating systems, certain portions of protected memory need to be redirected prior to allowing a user access to the system. In particular, in operating systems whose file systems store data on persistent storage, it is necessary to redirect the file system data itself to properly protect these systems. For example, in the Microsoft DOS and Windows™ operating systems, there are a number of different file systems, which store data on a system hard drive. These systems use a File Allocation Table (FAT) to track and allocate clusters to files. In particular, FAT file systems use a table of “next” cluster numbers. To work correctly in these environments, the DARS is implemented to handle both the FAT 16 system and the FAT 32 system. In FAT 16, the FAT is a table of 16 bit cluster numbers. In FAT 32, the table consists of 32 bit cluster numbers.

Figure 11 is an example block diagram of a disk drive with a File Allocation Table arranged according to the FAT 16 architecture. In Figure 11, a logical drive 1101 is organized as: (1) Boot record sector 1102, (2) FAT sector 1103, (3) Root directory sector 1104, and (4) Data clusters 1105. The first three sections 1102, 1103, and 1104 are fixed in size. The data section starts at a fixed point on the drive and is viewed by the operating system as a sequence of clusters.

Figure 12 is an example block diagram of a disk drive with a File Allocation Table arranged according to the FAT 32 architecture. A logical drive 1201 is organized as: (1) Boot record sector 1202, (2) FAT sector 1203, and (3) Data clusters 1204. The first two sections 1202 and 1203 are fixed in size. The data section starts at a fixed location and is viewed by the operating system as a sequence of clusters.

Figure 13 is an example block diagram of a disk drive arranged according to the New Technology File System (“NTFS”) architecture. The entire logical drive is organized using cluster segmentation, and all data is organized as files within the cluster framework. For example, cluster 1301 contains the Boot Sector.

The DARS is implemented to redirect all portions of the disk drive, including those used by the operating system, as shown in Figures 11, 12, and 13. This

feature allows a random hardware “power off” to leave the system secure. One method of implementing the DARS in FAT 16 and FAT 32 is to define two sets of the AST and PSRT redirection tables, which are organized to different data size abstractions. According to this embodiment, the redirection tables are sector-based for that part of the disk drive that precedes the data section and cluster-based for the data part (the remainder) of the disk drive. Because all available space on a drive will be in the data section of the drive, sectors from the pre-data part of the drive are redirected to the cluster organized part of the drive. Thus, several computations are necessary in this embodiment to translate sectors to clusters and vice versa.

To determine in which cluster a logical sector lies, the DARS in this example embodiment uses the following formula:

$$(1) \quad \text{Cluster} = ((\text{LSector} - \text{D0}) / \text{SPC}) + 2;$$

where,

D0 = logical sector number where the data section starts

SPC = # sectors per cluster.

The +2 results from a specific nuance in that the FAT 16/32 architecture labels the first cluster in the “data” section as Cluster 2.

To calculate the first sector of a cluster, the DARS uses the formula:

$$(2) \quad \text{LSector} = ((\text{Cluster} - 2) * \text{SPC}) + \text{D0}.$$

To calculate the offset within a cluster for a given logical sector, the DARS uses the formula:

$$(3) \quad \text{Offset} = (\text{LSector} - \text{D0}) \% \text{SPC}.$$

Where the “%” is a modulus operator that gives the integer remainder from a division.

Using the Microsoft New Technology File System “NTFS”, the operating system views the entire disk drive as a sequence of clusters, and all space on the drive is allocated in clusters. (See Figure 13.) Conversion from clusters to sectors or sectors to clusters is required because the file system usually handles cluster numbers, but reads and writes are usually handled as sectors. For example, queries for available space will usually

be obtained in terms of clusters. Therefore, the following example formulas can be used to translate between logical sectors and clusters:

$$(4) \quad \text{LSector} = \text{Cluster} * \text{SPC}$$

$$(5) \quad \text{Cluster} = \text{LSector} / \text{SPC}$$

For purposes of managing the redirection information generically, the DARS can be implemented using a concept of virtual clusters. In this implementation, the DARS only deals with redirection at the cluster level and is designed to be independent of the file system and the operating system.

Real cluster numbers are translated to virtual clusters using the formula:

$$(6) \quad \text{VCluster} = \text{RCluster} + \text{Offset};$$

where,

$$\text{Offset} = (\text{C0} / \text{SPC}) + \text{A};$$

C0 = the logical sector number of Real Cluster zero.

SPC = the number of sectors per cluster

$\text{A} = (\text{C0} \% \text{SPC}) ? 1 : 0$ ; Where the “?” is an operator, indicating that if the expression  $(\text{C0} \% \text{SPC})$  evaluates to “true” (any non-zero value), then  $\text{A} = 1$ . If the expression  $(\text{C0} \% \text{SPC})$  evaluates to “false” (zero), then  $\text{A} = 0$ .

The virtual cluster number that contains a logical sector is given by:

$$(7) \quad \text{VCluster} = (\text{LSector} - \text{SO}) / \text{SPC}$$

where,

$\text{SO} = (\text{C0} \% \text{SPC})$  is the offset of sector zero within Virtual cluster zero

The first sector of a VCluster is:

$$(8) \quad \text{LSector} = (\text{VCluster} * \text{SPC}) + \text{SO};$$

To calculate the offset within a virtual cluster for a given logical sector, the

DARS uses the formula:

$$(9) \quad \text{Offset} = (\text{LSector} - \text{SO}) \% \text{SPC}$$

In an embodiment that uses virtual clusters, the virtual clusters don't have to be redirected to physical clusters on the same drive. Clusters may be redirected to any place where storage is available, for example, on other drives, on network drives etc. The virtual cluster concept only requires an object of "available space" to which clusters may be redirected, and thus has the advantage of flexibility for a variety of operating systems. Virtual clusters may also be implemented to include virtual sectors or any other "virtual" data abstraction level. In systems that support object-oriented programming, the virtual cluster can be implemented as a class, whose implementation is specified by a particular instance of the Redirection Driver.

In some embodiments, a protection level is supported that allows the specification of some unprotected space. (For example, unprotected space 1004 in Figure 10.) Unprotected space can be implemented by allowing some set of sectors to be read and written into their original locations. Storage that is designated as unprotected is persistent through a reboot. By making the unprotected space appear to the operating system as a drive (which is designated to not be redirected), files can be stored in this space that will be persistent through a reboot. Specifically, because the Redirection Driver can be implemented to intercept reads and writes at the file level, files thus can be redirected into unprotected space – a special drive – instead of redirected space. The information in these files will then be persistent through a reboot. Redirection at the file level may be implemented such that the file appears to the user as if it is still in its original location.

In an alternative embodiment, an option is presented to a user to save the modifications attempted to protected space. For example, an option can be presented upon computer system shut down to save redirected data in this fashion. The DARS tables contain the information necessary to identify all modifications to the protected space. Thus, by transferring the information to the appropriate file allocation tables, the DARS can make the changes that have been made in the redirected space permanent in the protected space. For example, a user may wish to protect the data on a computer system while running a new software application in order to prevent unexpected data modifications. After running the software application and verifying that no unwanted

modifications were made, the changes that were made in the redirected space can be made permanent in the protected space. If unexpected data modifications were made, the user could simply re-boot the computer system, restoring the original data.

5 In an example embodiment of the DARS, implemented on a computer system running the Microsoft Windows 9X operating system, the Redirection Driver is implemented with the standard driver entry points that allow the operating system to communicate with any driver. The operating system supports the implementation of a Virtual Device Driver (a VxD), which can be inserted into the driver chain, so that the standard entry points of the driver are invoked by the operating system when the associated  
10 event is triggered.

Figures 14-18 are example flow diagrams describing the steps performed by the DARS Redirection Driver at each of five standard driver entry points. One skilled in the art will recognize that a Redirection Driver may be implemented with additional entry points or other entry points, and that the five entry points described with relation to Figures  
15 14-18 are intended only to be an example of a preferred embodiment, and are not intended to limit the scope of the invention.

Figure 14 is an example flow diagram of the steps performed by the Disk Access Redirection System when the computer system boots up. In step 1401, the computer system or a portion of the DARS determines the protection level for the computer  
20 system, as described with reference to Figure 4. In step 1402, the DARS examines the determined protection level, evaluating whether or not any protection is to be implemented for the computer system. If no protection is to be performed for the computer system, the Redirection Driver is not registered (not inserted into the driver chain) and will not intercept any disk access requests. If full or partial protection is designated, then the  
25 DARS continues in step 1405, to create an Available Space Table to map all of the available space for the appropriate disk drives. In step 1406, the DARS creates an Unprotected Space Table in embodiments that support partial protection. In step 1407, the DARS creates a Protected Space Redirection Table. In step 1408, the DARS maps the available storage space in the Available Space Table to initialize the AST. In step 1409,

the DARS maps unprotected space in the Unprotected Space Table to initialize the UST. For each storage device that is part of the computer system, steps 1408 and 1409 are performed. In one embodiment, a single set of DARS tables is used. In an alternative embodiment, a separate set of DARS tables is provided for each storage device in the system. Available space may be determined by any of a number of methods, for example, by querying the appropriate level device driver, by reading the FAT tables directly, by consulting a file containing addresses of such areas, accessing an access control list maintained by the operating system, or querying the user. The addresses that correspond to the unprotected space and protected space are determined from the input received after determining the desired protection level (step 1401). In step 1410, the system registers the Redirection Driver with the operating system as required by that system. For example, in systems running a Windows 9x operating system, any driver stored in the appropriate directory is loaded and put in the correct location to intercept access requests when the system boots. In systems running other operating systems, a different registration process may be required.

In an alternate embodiment of the invention, a user may specify a computer system protection level that allows modifications to be maintained across several system reboots. In step 1403, if the determined protection level indicates that the Available Space Table, Protected Space Redirection Table, and, optionally, the Unprotected Space Table are to be retrieved from a previous session, then the DARS proceeds to step 1404. In step 1404, the DARS loads the existing tables, preferably from persistent memory into volatile memory, and then proceeds to step 1410 where the Redirection Driver is registered to intercept disk access requests, as described above.

Figure 15 is an example flow diagram of the steps performed by the Disk Access Redirection System when a new device is mounted by the computer system. In step 1501, the DARS determines whether or not the newly mounted device is a storage device, and, if not, it terminates, else it continues in step 1502. Steps 1502 and 1503 are similar to steps 1408 and 1409, respectively, described above. In step 1502, the DARS maps the available space in the Available Space Table. In step 1503, the DARS maps any



designated unprotected space in the Unprotected Space Table. In one embodiment, a user interface may need to be presented to query the user as to whether it is desired to protect this new device (for example, when partial protection has been indicated). In step 1504, when the operating system so requires, the Redirection Driver is registered to intercept read and write requests to the new storage device.

Figure 16 is an example flow diagram of the steps performed by the Redirection Driver when a read request is intercepted. In step 1601, the driver determines whether the requested address is in the UST, indicating unprotected space that may be freely written to and read from. If so, then the driver proceeds to step 1604 to read from the requested address, else, it proceeds to step 1602. In step 1602, the driver determines whether or not the requested address is mapped in the PSRT, indicating that the protected space already has been redirected and the redirected address needs to be read from instead of the original address. If the requested address has been redirected, then the driver continues in step 1603, else, it continues in step 1604. In step 1603, the driver sets the address to the redirected address already mapped in the PSRT. In step 1604, the system forwards the read request, with either the original or redirected address, to the next I/O driver in the driver chain to perform the read.

Figure 17 is an example flow diagram of the steps performed by the Redirection Driver when a write request is intercepted. In step 1701, the driver determines whether or not the requested address is in the UST, indicating unprotected space that may be freely written to. If so, then the driver continues in step 1707 to write to the requested address, else it continues in step 1702. In step 1702, the driver determines whether or not the requested address is mapped in the PSRT, indicating that the protected space already has been redirected. If the requested address has not yet been redirected, then the driver continues in step 1703, else it continues in step 1706 to retrieve the redirected address. In step 1703, the driver identifies and allocates available space from the AST as required to satisfy the request. The original write request is progressively broken up into smaller requests, if a large enough space is unavailable to satisfy the request. In step 1704, the driver writes a new record to the PSRT, mapping the requested protected address to the

available address allocated from the AST in step 1703. In step 1705, the driver removes the allocated space from the AST, because, once the available space is used as redirected space, it is no longer available. In step 1706, the driver sets the address requested in the disk access request to the redirected address as mapped in the PSRT. In step 1707, the driver forwards the write request with the original unprotected or redirected address to a lower level disk driver to perform the write.

As described above, in an alternate embodiment of the DARS, protection levels are implemented that allow modifications to the redirected space to remain persistent across one or more system reboots. In one such embodiment, the standard driver entry point for system shutdown is implemented in the Redirection Driver. Figure 18 is an example flow diagram of the steps performed by the Redirection Driver when a system shutdown request is intercepted. In step 1801, the driver determines whether or not the protection level in effect for the computer system indicates saving the DARS tables. If so, then the driver continues in step 1802, else, the on\_shutdown routine terminates. In step 1802, the driver writes the Available Space Table to persistent storage. In step 1803, the driver writes the Unprotected Space Table to persistent storage. In step 1804, the driver writes the Protected Space Redirection Table to persistent storage. Persistent storage for these purposes may also be a logical drive in the file system that has been designated as unprotected storage, as described above. One skilled in the art will recognize that other steps may be performed on computer system shutdown in this routine, for example, to store unprotected areas if these had been redirected to redirected space instead of written to directly by the appropriate driver.

Although specific embodiments of, and examples for, the present invention are described herein for illustrative purposes, it is not intended that the invention be limited to these embodiments. Equivalent methods, structures, processes, steps, and other modifications within the spirit of the invention fall within the scope of the invention. For example, the teachings provided herein of the present invention can be applied to any systems with associated persistent data storage, for example, a personal computer system

with a hard disk drive, or a networked server system with remote data repositories. In addition, the teachings may be applied to other types of systems where driver-like code may be implemented to redirect the flow of data through the system. These and other changes may be made to the invention in light of the above-detailed description.

- 5 Accordingly, the invention is not limited by the disclosure, but instead, the scope of the present invention is to be determined by the following claims.